

Connectionist Modeling of Developmental Disorders

Annette Karmiloff-Smith

Neurocognitive Development Unit,
Institute of Child Health, London.

Michael S. C. Thomas

Neurocognitive Development Unit,
Institute of Child Health, London.

Running head: Developmental disorders

Address for correspondence:

Annette Karmiloff-Smith or Michael Thomas

Neurocognitive Development Unit,

Institute of Child Health,

30, Guilford Street,

London WC1N 1EH, UK.

Email: M.Thomas@ich.ucl.ac.uk

Tel.: 020 7905 2747

Fax: 020 7242 7177

Introduction

Connectionist models have recently provided a concrete computational platform from which to explore how different initial constraints in the cognitive system can interact with an environment to generate the behaviors we find in normal development (Elman et al., 1996; Thomas & Karmiloff-Smith, 2002a). In this sense, networks embody several principles inherent to Piagetian theory, the major developmental theory of the twentieth century. By extension, these models provide the opportunity to explore how shifts in these initial constraints (or boundary conditions) can result in the emergence of the abnormal behaviors we find in atypical development. Although this field is very new, connectionist models have already been put forward to explain disordered language development in Specific Language Impairment (Hoeffner & McClelland, 1993), Williams Syndrome (Thomas & Karmiloff-Smith, 2002b), and developmental dyslexia (Seidenberg and colleagues, see e.g. Harm & Seidenberg, in press); and to explain unusual characteristics of perceptual discrimination in autism (Cohen, 1998; Gustafsson, 1997). In this entry, we examine the types of initial computational constraints that connectionist modelers typically build in to their models, and how variations in these constraints have been proposed as possible accounts of the causes of particular developmental disorders. In particular, we examine the claim that these constraints are candidates for what will constitute innate knowledge. First, however, we need to consider a current debate concerning whether developmental disorders are a useful tool to explore the (possibly innate) structure of the normal cognitive system. We will find that connectionist approaches are much more consistent with one side of this debate than the other.

Developmental disorders and modularity

Cognitive neuropsychology assumes that the adult cognitive system has a modular structure, whereby the system can be decomposed into specialized functional components (although whether these components correspond to localized areas of the brain, or can be captured by brain imaging techniques remains an open question). In addition, cognitive neuropsychology assumes that that selective behavioral deficits in adults with brain damage can reveal this modular structure. Developmental disorders can also produce apparently specific deficits in the endstate of development. For example, Williams syndrome (WS), a developmental disorder caused by a microdeletion of contiguous genes on one of the alleles of chromosome 7, is characterized by a behavioral profile of relative proficiency in language, face processing, and theory-of-mind (attributing mental states to others), but severe deficits in other skills such as visuospatial processing, number, and problem solving (Karmiloff-Smith, 1998). In hydrocephalus with associated myelomeningocele (a protrusion of the membranes of the brain or spinal cord through a defect in the skull or spinal column), language can be the only area of proficiency. Individuals suffering from Specific Language Impairment (SLI) show the opposite pattern, often apparently performing within the normal range in all domains except language. In autism, even individuals with normal IQs are selectively impaired in tasks that require judging another's mental states (Baron-Cohen, Tager-Flusberg, & Cohen, 1993).

Evidence in many of these disorders of specific high-level deficits at the end of development has encouraged some researchers to view developmental disorders as

offering the same theoretical insights into the static structure of the cognitive system that cases of adult brain damage provide. Where such disorders are of a genetic origin selective developmental deficits are interpreted as revealing innate underpinnings for such structure. For example, Baron-Cohen et al. (1993) have argued that, in individuals with autism, an apparent deficit in reasoning about mental states can be explained by the impairment of an innate, dedicated module for such reasoning (the 'Theory of Mind' module). Van der Lely (1997) maintains that selective behavioral deficits in the language performance of children with grammatical SLI can be explained by damage to an underlying, innate module representing syntactic (rule-based) information. Clahsen and Almazan (1998) have proposed that a behavioral deficit in WS language supports the view that while their syntactic skills are 'intact', individuals have a deficit to a specific aspect of their (modular) language knowledge, that of accessing information about words that are exceptions to syntactic rules.

However, there are a number of problems with the adult brain damage approach to developmental disorders (Karmiloff-Smith, 1998). These boil down to the suspicion that such an approach massively underestimates the complexity of the path from gene to behavior. The prevalence of many-to-many, very indirect mappings in the relation of genes to cognition undermines the claim that direct specific mappings will exist between particular genes and individual high-level cognitive abilities. To the extent that genes are involved in the causal chain of several cognitive domains, it will be less likely that they code anything specific to a single domain. Indeed, direct mappings are unlikely given that spatial distributions of gene expression in the brain are rarely narrowly confined to subsequent areas of functional specialization in the adult brain and therefore seem unable

to code directly for domain-specific developmental outcomes. Even if they could, the idea that behavioral deficits identified in the endstate of a developmental disorder could reflect the impairment of a single functional module is predicated on the dubious assumption that the rest of the cognitive system could nevertheless develop normally. For this to be true requires either that functional modules develop independently of overall brain growth, or that the content of modules is fixed in advance (i.e. the content is innately specified). But neither of these assumptions is likely to be true. With regard to the first, Bishop (1997) has argued persuasively that interactivity rather than independence is the hallmark of early development. With regard to the second, it seems likely that modular structure in the cognitive system and in the brain is an outcome of development rather than a precursor to it, and that the neonate brain does not support innate representations with specific content (Elman et al., 1996; Johnson, 1997). A growing number of studies show how both neural localization and neural specialization for biologically important functions, such as species recognition (Johnson, 1997) and language (Neville, 1991), take place gradually across development.

An alternative to the use of the adult brain damage model is the neuroconstructivist approach (Elman et al., 1996; Karmiloff-Smith, 1998) which views developmental disorders in terms of different developmental trajectories, caused by initial differences at a neurocomputational level. Thus there might be differences in the microcircuitry of the brain or the firing properties of neurons, as opposed to discrete lesions to particular large-scale brain structures or pathways. In this view, development is an interactive process in which the cognitive system self-organizes in response to interactions with a structured environment. Interestingly, this approach suggests that people with developmental

disorders may exhibit strengths as well as weaknesses. This prediction is consistent with superior face recognition skills found in WS and superior perceptual discrimination abilities found in autism. Neuroconstructivism further suggests that equivalent behavior across normal and abnormal phenotypes may mask different underlying cognitive processes. The notion that an ability is 'intact' or 'spared' because there is no apparent deficit at the behavioral level employs terminology from the adult brain damage model that may be misleading. To take an example, people with Williams syndrome can display scores on some language and face processing tasks which are in the normal range. Nevertheless, closer examination suggests different cognitive processes underlie the equivalent behavioral scores (Karmiloff-Smith, 1998).

Initial constraints in connectionist models

Current connectionist models of developmental disorders employ alterations to the initial model constraints which, after training, lead to an endstate exhibiting behavioral deficits. Connection weights are usually randomized so that the normal network has no specific knowledge prior to training. Thus it follows that the atypical network has no specific deficit in knowledge prior to training. The behavioral deficits that emerge when these atypical networks are trained are often quite different to the effects of damaging a normal network model after training has been completed. This holds even when the network manipulations are the same in each case. Thus current connectionist models are more consistent with the neuroconstructivist approach to developmental disorders than the adult brain damage approach.

What, then, are the initial constraints that connectionist modelers build into their

models of cognitive development? And how do changes to these constraints alter the trajectory of development? In fact, the constraints that connectionist models build in are quite strong ones, and this may come as a surprise to some readers. Connectionist models are often mis-characterised as being unitary / homogeneous / seamless / undifferentiated domain-general learning devices, whereby the environment is all-powerful in shaping the behavior of the final system. In fact, current connectionist models have a great deal of pre-existing structure built into them prior to any exposure to their training environment. What is general about connectionism are the principles of computation (Seidenberg, 1993). When the general principles of computation are combined with the boundary conditions for a specific domain, the result is a domain-specific model. It is the generality of these principles that gives the connectionist approach its explanatory power. That is, connectionism seeks not just to formulate descriptive generalizations about empirical phenomena, but to show how they derive from underlying and independently motivated principles (Seidenberg, 1993). However, connectionist models are just as reliant on the constraints of a given domain as they are on the computational principles. Without justified limitations in the design of network models, they become overpowerful data-fitting devices that can at best provide descriptively adequate accounts of cognitive abilities. In short, connectionist models of development do include initial structure but not initial representational content (Elman et al., 1996). The point is that in interacting with a training environment, networks create representational content and become increasingly structured. This additional structure reflects the nature of the training environment.

The structure or boundary conditions that these models build in prior to training

typically involves the following:

- (1) The initial state of the network, in terms of the number of units, layers, connections, and the pattern of connectivity, collectively known as the network architecture. The architecture determines the computational power of the network, and the type of computations to which the network will be suited. For example, recurrent networks are suited to processing sequential information, whereas associative networks are suited to pattern recognition. The a priori choice of the architecture will have a central role in determining the adequacy of the network in modeling a given domain of cognitive development. A reduction in the number of processing units, or the elimination of internal (hidden) processing units, will restrict the computational power of the network and, depending on the nature of the domain, cause some or all parts of the problem to be learned inadequately. Addition of layers of internal units beyond a single layer tends to delay learning, without a marked increase in effective computational power. Increase of hidden units within a single layer tends to improve performance on the training set, but may impair generalization beyond the training set.

- (2) The way a particular cognitive problem is presented to the network, in terms of the input and output representations. For a given domain, the representations determine the nature of the computational problem that the network will face. When the network has to extract a function from the training set (such as a general rule or regularity), the representational scheme will be crucial in determining how transparent or opaque this

function is to the network. For instance, if a network is given a training set in a form that masks the relevant similarity between those items in the problem domain that obey a rule, the network will have difficulty in extracting this regularity.

- (3) The learning algorithm that the network will use to change its connection weights (and potentially, its architecture). Most networks are trained by changing weights to minimize some cost function, such as the difference between the actual output and a target output. The rate at which weights are changed can have an impact on the success of a network in learning a problem. In particular, in complex domains, if weights are changed too quickly, the network may commit too early to a partial solution to the problem, and be resistant to change with subsequent training. The learning algorithm is key in determining the plasticity of the network to further learning. Some algorithms allow on-line changes to network architecture depending on how well the network is learning a problem. The way that the network's computational power is altered on-line will again have a considerable influence on the success of the network in capturing a cognitive ability (see (1) above).
- (4) The regime of training which the network will undergo. After modelers determine the network and the learning rule, they then expose the network to a training set. Often the network is exposed to the entire training corpus from the start of training. However, in some cases, the network might be trained on an initially limited training set, perhaps based on assumptions about the nature of a child's early learning environment. This initial restriction will affect later training. It may aid learning if the

smaller set is representative of the larger set, or if it allows the construction of internal representations that will be useful in learning the larger set. On the other hand, it may impair learning, if the initial training set contains detail irrelevant to the full domain. Alterations in network parameters early in training may have the same effect as restricting the initial training set, as Elman's work on learning syntax with recurrent networks has demonstrated (Elman et al., 1996).

In the connectionist framework, these constraints represent some of the candidates for innateness. (Although, equally, in principle any of the constraints could also be experience dependent). Alterations in one or more of these constraints may then lead to the emergence of disordered representations and impaired behavior in a model of atypical development.

We have noted that these models do not support innate representational content in that their weights are initially randomized. However, it is an open question about whether initial computational constraints (along with sensory input determined by the individual's interaction with the environment) are sufficient to drive development. One possible addition is the postulation of innate attentional predispositions (Elman et al., 1996). In this theory, innate knowledge is built into the subcortical part of brain in the form of a low acuity predisposition to attend certain stimuli. This predisposition then ensures the representation of input that will subsequently drive learning in the more powerful cortical system. For example, such an innate predisposition in face recognition might encourage the newborn infant to attend preferentially to visual stimuli containing a single blob positioned centrally below two blobs (see Johnson, 1997). Innate predispositions provide

another candidate factor that might be altered in the start state of the atypical system.

In general, current connectionist models of normal development do not restrict themselves to computational constraints and innate attentional predispositions in their start states, since these models incorporate high-level, domain-specific representations. These models must therefore be seen as a halfway house. In future they need to be extended to show how their domain-specific representations may emerge from some prior process operating over lower-level information (and with its own computational constraints).

We now turn to consider recent examples of connectionist models of developmental disorders.

Recent models

Autism. Autism is a developmental disorder characterized primarily by a central triad of deficits in social interaction, communication, and imagination, but also by a range of secondary deficits. These include a restricted repertoire of interests, an obsessive desire for sameness, excellent rote memory, improved perceptual discrimination, and an impaired ability to form abstractions or generalize knowledge to new situations. Evidence from neuropathological investigations of the brains of affected individuals is suggestive of abnormal wiring patterns in various brain regions, perhaps caused by deficits in neuronal migration during fetal development, curtailment of normal neuronal growth, and/or aberrant development. Cohen (1998) has argued that these structural deficits are consistent with too few neurons in some brain areas, such as the cerebellum, and too many neurons in other areas, such as the amygdala and hippocampus. Cohen has

proposed that simple connectionist models trained on categorization tasks can link such differences in neurocomputational constraints to some of the secondary deficits found in autism. In some cases, children with autism have trouble acquiring simple discriminations and attend to a restricted range of stimuli, while in others, children have good discrimination and indeed very good memory but seem to rely on representing too many unique details of stimuli. Cohen showed that simple backpropagation networks with too few hidden units showed a failure to learn classifications tasks, while those with a surfeit of hidden units showed very fast learning, but subsequently generalization became poor, and the network increasingly responded according to particular details of the training set.

In a related proposal, Gustafsson (1997) speculated that the combination of a failure to generalize and heightened perceptual discrimination might be traced to the atypical development of cortical feature maps. In particular, he suggested that higher-than-normal levels of within-layer inhibition in the initial cortical sheet would lead to overly fine-tuned perceptual features, permitting good discrimination but preventing good generalization.

Although Gustafsson did not support this idea with implementations, related work by Oliver et al. (2000) offers an insight into how feature map formation can be developmentally disrupted by changes in the initial properties of a self-organizing connectionist network. Oliver et al. employed a neurobiologically constrained network in which a two-dimensional output layer received information from a single input retina. The network was shown a set of stimuli in the form of bars lying across the input retina. Oliver et al. showed that, using their initial parameter set, the output layer formed a topographic map of the possible inputs: certain areas of the output layer specialized in responding to each

input, and areas representing similar inputs were adjacent to each other in the output layer. Oliver et al. then re-ran the model, disrupting the network in different ways prior to exposing it to the training stimuli. They varied the threshold of the output units, disrupted the connectivity between the input and output layers, disrupted the connectivity responsible for lateral inhibition in output layer, and changed the similarity of the input stimuli to each other. Importantly, these manipulations demonstrated that tiny differences in the initial constraints under which the model developed could have a very significant impact on the outcome of development. The resulting topographic map suffered a range of disruptions, including output units failing to specialize at all or simply turning off, specialization emerging but not in organized areas, and organized areas emerging but without adjacent areas representing similar looking bars.

Much work remains to be done to develop these proposals exploring the neuro-computational underpinnings of developmental disorders such as autism, and in particular, to relate disordered feature maps or hidden unit numbers to higher-level cognitive deficits such as those in social interaction, communication, and imagination that characterize autism. Nevertheless, such work importantly illustrates a new conception of such disorders in terms of development operating under atypical constraints, rather than in terms of deficits to high-level modules in a static model of the normal adult cognitive system.

Developmental dyslexia. This disorder has been the focus of much connectionist research, given the success of models in capturing the normal processes of reading. A number of models by Seidenberg and colleagues have sought to change initial constraints in models of reading to simulate either surface dyslexia (where the subject has difficulty

reading words which are exceptions to normal rules of pronunciation), phonological dyslexia (where the reading of novel words is impaired), or a combination of both. Typically these models learn mappings between codes representing orthography, phonology, and semantics. Surface dyslexia has been simulated by employing ‘too few’ hidden units in the model, or by reducing the learning rate. Phonological dyslexia has been simulated by degrading the phonological representations in some way, for instance in the type of coding scheme used. For example, Harm and Seidenberg (1999) pre-trained one part of their model to develop appropriate phonological representations, prior to learning the reading task. When this ‘phonological’ part of the model was impaired, either by reducing its initial computational power or by limiting the size of the connection weights it could develop, the result was a network exhibiting phonological dyslexia at the end of training on the reading task.

Specific Language Impairment. Hoeffner and McClelland (1993) sought to capture deficits found in the morphosyntax of subjects with SLI, specifically their difficulty in the learning of rule-based inflectional morphology in verbs. Hoeffner and McClelland employed an attractor network mapping between semantic codes and phonological codes. They simulated SLI by changing the initial phonological representations, in line with a hypothesis that SLI may be caused by early perceptual impairments. Specifically, they impaired the network’s ability to represent word-final stops and fricatives (including /t/, /d/, and /s/). Although the model they used didn’t show an ideal fit to the normal data when unimpaired, it nevertheless captured a number of the key deficits of SLI when trained with impaired representations, particularly a selective difficulty with the

formation of regular (+ed) past tenses. In this case, the initial phonological deficit obscured precisely the information that the network required to be able to learn the relevant generalizations about regular past tense formation. However, it should be noted that the perceptual deficit account of SLI remains controversial, and this disorder may well turn out to be heterogeneous, with several different causes.

Williams Syndrome. Recent work in our laboratory has examined underlying deficits in the language of individuals with Williams Syndrome (WS). Initially thought to be ‘spared’, closer examination revealed a number of subtle deficits in their language. It had been reported that individuals with WS show difficulties in forming the past tense of irregular verbs, while showing good performance on the regular, rule-based past tense formations (Clahsen & Almazan, 1998). Our recent empirical work suggests that much of this apparently selective deficit is due to an overall delay in language development (young children also find irregular verbs difficult). However, individuals with WS do appear to show a deficit in generalizing knowledge of inflectional patterns to novel forms. These two patterns (selective difficulty with irregular inflections vs. reduced levels of generalization) continue to be argued for in the literature on WS language.

Thomas and Karmiloff-Smith (2002b) set out to explore whether alterations to the initial computational constraints of a connectionist model of past tense development could account for either of these patterns of data. The past tense network mapped from verb stem to past tense form in the presence of lexical-semantic information. Various theoretical claims have been made that the WS language system develops under different constraints. These include the proposals that their phonological representations may be

atypical and perhaps rely on sensitive auditory processing; that their lexical-semantic representations may show anomalous organization; or that lexical-semantic information about words may be poorly integrated with phonology. In order to explore the viability of these different accounts to explain the pattern of performance in the past tense task, Thomas and Karmiloff-Smith altered the initial constraints of the network model to implement each type of lower-level deficit.

The results revealed that reduced generalization was consistent with atypical phonological representations (specifically with reduced similarity and redundancy), or interference in integrating lexical-semantic and phonological knowledge. A range of computational constraints caused poorer acquisition of irregular verbs – these verbs are in the minority and their acquisition can be disrupted under non-optimal learning conditions. However, attenuated activation from lexical-semantics was able to selectively delay irregular acquisition, offering a link to one of the proposed deficits in WS. Finally, the model demonstrated for the first time precisely how different computational constraints interact in a system in the process of development: the atypical trajectory found in WS may arise from the combination of more than one altered constraint.

Further empirical work remains to be carried out to clarify which of the WS patterns is the correct one. However, this modeling work has determined which of the competing accounts are viable within an existing developmental framework, and therefore provided a focus for future investigations. Once again, this contrasts with previous theoretical work that has construed the WS language system in terms of selective deficits to a static model of the normal adult language system (e.g., Clahsen & Almazan, 1998).

Discussion

Developmental disorders can inform the study of normal development because they provide a broader view of the parameter space within which development takes place. The empiricist argues that the environment specifies a capacity so strongly that systems with a wide variety of initial structures must come to reflect it. The nativist argues that the environment specifies a capacity so poorly that the system must come equipped with pre-existing structure if it is to find the correct solution given the environmental input. The neuroconstructivist argues that the robustness of the cognitive system to changes in its initial set-up (so long as we can come to understand precisely what these changes are) will reveal much about how evolution has placed its bets concerning the capacities that can be trusted to emerge through experience, and the capacities which must be given a firmer guiding hand through development. Connectionist models provide a powerful tool with which to investigate the role of initial computational constraints in determining the trajectory of both typical and atypical development. For developmental disorders in which selective deficits in high-level behaviors are reported, the connectionist framework ensures that these deficits are properly seen in terms of the outcome of the developmental process itself.

References

Baron-Cohen, S., H. Tager-Flusberg, and D. J. Cohen, 1993. Understanding Other Minds: Perspectives from Autism, Oxford: Oxford University Press.

Bishop, D. V. M. 1997. Cognitive neuropsychology and developmental disorders: Uncomfortable bedfellows, Quarterly Journal of Experimental Psychology, 50A:899-923.

Clahsen, H. and M. Almazan, 1998. Syntax and morphology in Williams Syndrome, Cognition, 68:167-198.

Cohen, I. L., 1998, Neural network analysis of learning in autism, in Neural networks and psychopathology, (D. Stein & J. Ludick, Eds.), Cambridge: Cambridge University Press, pp. 274-315.

* Elman, J. L., E. A. Bates, M. H. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett, 1996, Rethinking innateness: A connectionist perspective on development, Cambridge (Mass): MIT Press.

Gustafsson, L. 1997. Inadequate cortical feature maps: A neural circuit theory of autism, Biological Psychiatry, 42:1138-1147.

Harm, M. & Seidenberg, M. S, 1999, Phonology, reading acquisition, and dyslexia: Insights from connectionist models. Psychological Review, 106:491-528.

Hoeffner, J. H. and J. L. McClelland, 1993, Can a perceptual processing deficit explain the impairment of inflectional morphology in developmental dysphasia? A computa-

tional investigation, in Proceedings of the 25th Child language research forum, (E. V. Clark, ed.), Stanford University, CA: Center for the Study of Language and Information, pp. 38-49.

* Johnson, M. H., 1997, Developmental Cognitive Neuroscience, Oxford: Blackwells Publishers.

* Karmiloff-Smith, A., 1998, Development itself is the key to understanding developmental disorders, Trends in Cognitive Sciences, 2:389-398.

Neville, H. J. 1991. Neurobiology of cognitive and language processing: Effects of early experience, in Brain maturation and cognitive development: Comparative and cross-cultural perspectives. Foundation of human behavior, (K. R. Gibson and A. C. Peterson, eds.), New York: Aldine de Gruyter, pp. 355-380.

Oliver, A., Johnson, M. H., Karmiloff-Smith, A., & Pennington, B., 2000, Deviations in the emergence of representations: A neuroconstructivist framework for analysing developmental disorders, Developmental Science, 3:1-23.

Seidenberg, M. 1993. Connectionist models and cognitive theory, Psychological Science, 4:228-235.

* Thomas, M. S. C. & Karmiloff-Smith, A, 2002a, Modeling typical and atypical cognitive development, in Handbook of Childhood Development, (U. Goswami, Ed.),

Oxford: Blackwells Publishers, pp. 575-599.

Thomas, M. S. C. & Karmiloff-Smith, A, 2002b, Modeling language acquisition in atypical phenotypes. Manuscript submitted for publication.

Van der Lely, H. K. J., 1997, Language and cognitive development in a grammatical SLI boy: Modularity and innateness, Journal of Neurolinguistics, 10:75-107.